



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Report on the performance of a large-area, gamma-ray imager for search

L. Fabris, K-P Ziock

October 19, 2005

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This work was performed under the auspices of the U.S. Department of Energy by University of California, Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

Report on the performance of a large-area, gamma-ray imager for search

Lorenzo Fabris, Klaus-Peter Ziock

Introduction

We are currently constructing a prototype, large-area, gamma-ray detector for conducting vehicle-mounted, mobile-search operations. The system is unique in that it relies on imaging to discriminate point sources of interest from the natural background variations. In a non-imaging instrument the background fluctuations mimic the signature seen from real sources at a distance and one is limited in sensitivity to detecting only those sources that overwhelm the local background variations—not just the counting statistics associated with a given measurement. The net result is that a larger detector is generally not more sensitive to detecting sources in the world at large. [1, 2]

In a previous publication [3] we reported on the detection of a 1-mCi source at more than 80 meters from the detector using a proof-of-principle instrument (see Fig. 1) constructed to demonstrate how imaging removes the size limit on search instruments. In this report we document a systematic effort using the same detector to demonstrate that imaging detectors can reliably detect weak radiation sources at many 10's of meters. Specifically, we collected data on a 1-mCi ^{137}Cs source 65 m from the path of the search instrument.

Receiver Operator Curves

The key performance indicators for a search instrument are the probability of detection and the probability of false alarm. The former tells how likely it is that a source one really wishes to detect is missed. The latter tells one how often the system falsely claims the presence of a source when none is there. Obviously one wants a very high probability of detection while simultaneously maintaining a very small false alarm probability. A standard means of displaying these two properties simultaneously is to plot one as a function of the other in a Receiver Operator Curve (ROC). An excellent ROC would have no probability of false alarms with a 100% probability of detection and would look like the red plot in Fig. 1. Less-than-ideal performance would fall below and to the right of this line. A system that relies solely on guessing would have a linear ROC as shown by the blue line. In the real world, systems generally fall between these two ideal limits.

Methodology

The ideal means of generating an ROC is to take a series of measurements with and without the source and then calculate the number of false alarms from the no



Fig. 1. The prototype imager is mounted in the rear of a rental truck. The image is formed from the shadow pattern cast on the 3 x 19 array of detector elements by the lead mask.

source data and the number of detections from the source data. We are attempting to detect sources that contribute less radiation flux to that seen by the detector than natural variations in the background. This means that the issue is not can we detect a source of a given strength, but rather, can we see it without alarming on the natural structure in the background in the world at




Fig. 1. Receiver Operator Curves ROC for instruments ranging from ideal (red) through useless (blue.)

large. The best way to do this is to place the source at different locations in the world and see if one detects it. Unfortunately, it is difficult to leave radiation sources at many different locations due to administrative regulations and safety concerns. While this makes determining the probability of detection difficult, determining the false-alarm rate is relatively simple. One can move the detector around different regions and see how many times it alarms at a given detection threshold in the absence of a source.

To generate the probability of detection we have used a two part approach. Experimental data was collected on the type of source we are interested in detecting. A large number of passes past the same location are made with and without the source (see Fig. 3). The data is then replayed with the no-source runs subtracted from the source runs. This generates a high-fidelity model of the instrument response to the source only. It is independent of any structure in the background other than counting statistics. This model is then scaled to the correct source strength (thereby reducing the statistical noise) and injected into real data collected in the community around Livermore. This gives us the ability to synthetically put the source in many locations and with varying strengths, so that a probability of detection can be generated for the same areas as the false alarm data.

Model Generation

This approach is only possible because we collect list-mode data from the instrument. This means that all the information available during an acquisition about each event is recorded to disk. This allows us to replay the data, varying parameters as if we were taking another acquisition, i.e. we can change energy cuts, calibration files, location cuts, etc.

Because the Large Area Imager generates images using an indirect imaging approach, great care was taken to fully include the imaging effects in the model. For the details of how the data is handled, the interested reader is referred to [3]. In short, an image of the shadow mask on the detector is allowed to accumulate while the detector is within a single image pixel width. When a pixel boundary is reached, the image generation algorithm is performed and the shadow pattern is turned into an image of the current portion of the world in the field of view. This is then added to the world-view maintained by the system. The statistical errors associated with the data manipulations are known and tracked, allowing the significance of the counts in any image pixel to be known.

In developing the model, we were concerned that a faithful recreation of the response function was generated, i.e. that the result would be the same as if a source



Fig. 3. The source is placed inside the near wall of the building which is on the road ~ 65 m from the detector track. No effort is made to correct for the effect of obscuring objects between the imager and the detector.

were really present during an acquisition. To this end, the model does not save the world-map generated as a final product, rather we save the individual detector images of the shadow mask *before* the image reconstruction is applied. To add the model source to a background run, the scaled detector maps are added to the background run detector maps at each location before the data is deconvolved to obtain an image.

The model data was collected on two separate occasions. On each date we collected 5 full sets of data past the 1-mCi, ^{137}Cs source at a range of 65 m. Each full set comprises a mask and an anti-mask pass with and without the source present. The model then comprises 10 no-source runs subtracted from 10 source runs. Each run includes a mask and an anti-mask pass. When the model is made, we only need to keep that part of the track past the source that actually has the source in the field of view. This comprises approximately 133 m of distance along the track. Since the absolute source location in the image was not known exactly, the model keeps data along 150 m of the track. Because of the way world images are constructed this means that there are some contributions to the image over nearly 300 m of the final image.

The image of the full model at 65 meters is shown in Fig. 4. On the whole the results obtained using the model are excellent as can be seen by comparing the left and right plots in Fig. 5. On the left we show the result of the model added back into one of the no-source data sets used to generate the model. On the right we show a representative single source data set used to generate the model. As can be seen, the two are qualitatively identical. Note that there is slightly larger




Fig. 4. *Unscaled model. The source contributes to the data over 130 m of transit. The contribution to the image occurs over twice this distance.*

statistical noise due to the addition of the model in the left image. This is unavoidable since the model data has counting statistics errors associated with it and since there is no way to “smooth” or fit the results without the risk of imposing a non-physical effect into the model. This means that the results are somewhat less sensitive than if the source had really been placed in the world at large.

Data Collection

Whenever one is concerned with detecting sources there are effects beyond the control of the operator. In particular, one can always add enough shielding to a source to completely mask its radiation (in this case, gamma-ray) signature. This is a separate problem that falls outside the scope of this report. Given a shielding model, one can recast the problem for a given shielded source into one for a weaker unshielded source (with some shift in the energy spectrum). Our real interest is

to define the false alarm probability so that it is unambiguous. For the purposes of this report, we are interested only in false alarms from the fixed background and its variation as one passes through the environment. We are not interested in “nuisance” alarms that might occur from real sources that one is not worried about. (In principle these can be removed using energy resolution.) To this end we have opted to look for a ^{137}Cs source as the characteristic 662 keV emission from this source is sufficiently removed from the most common nuisance alarms (medical patients) that it is unlikely to be detected. Never-the-less, we did make one very strong detection of a source, obviously in transport on a highway neighboring our search path. This entire data set was excluded from the analysis.

To generate the ROC curves, we have two sets of data collected in the Livermore area. One comprises a total of 4.5 km of data taken at 16 km/h in 6 separate integrations. The other comprises a total of 10.6 km of data taken at 40 km/h in 9 separate integrations. Each integration includes a mask and an anti-mask pass past the same locations. This represents 16.8 and 15.9 minutes of search at the respective velocities. The sigma of detection for each of the pixels at 65 m is used to generate the probability of false alarm. There are a total of 644 and 1513 pixels (or locations) respectively in each analysis set.

To determine the probability of detection, the model data is added to each of these data sets. Because the length of the model is only 150 m, it can be added at many different locations using an offset from the start location to each of the longer background runs. To get as many locations as possible, and yet maintain statistical independence, the source is added at sequential locations shifted by 75 m. This results in a total of 42 and 137 locations, respectively, for the two different velocity data sets. The problem is rerun with the source at




Fig. 5. *An image obtained at 65 m with no source present but with the model added in to represent a 1 mCi source (left.) On the right is one of the 10 mask/anti-mask images with a source actually present.*




Fig. 6. ROC results for 16 km/h data. The source is at 65 m from the imager track. The performance for the 1.5 mCi source is perfect, (for the number of trials available,) indicating that the system can detect weak sources at significant distances.

different strengths. The ROC statistics generated are shown in Figs. 6 and 7.

Discussion

As can be seen, the imager does quite well at detecting sources as weak as 1 mCi. However, it should also be noted that these results are conservative. First, there is no shielding around the detector. Thus, the background is worse than it will be for the final deployable instrument. Second, as mentioned above, the use of the model approach adds a small but finite amount of counting statistics noise to the results. Third, in the current instrument, the data is displayed at the same resolution as the base imager resolution. This means that if a source happens to fall in between two pixels in the world view, it will show up split between the two, but with half of the peak intensity it would have if it fell entirely in a single pixel. To help reduce this in the final instrument, we will generate a world map at twice the resolution of achieved with the imager. This over-sampling technique helps reduce this aliasing effect by a factor of two. A final reason that the observed performance is a worst-case result originates from the fact that the data was collected in several separate passes. Because these were started manually by sighting on objects along the roadside, a several meter error can occur. This shifts the various images with respect to each other, resulting in a smaller peak. In the final imager the mask and the anti-mask data will be obtained in a single pass, removing this problem as well.




Fig. 7. ROC results for 40 km/h data. The source is 65 m from the imager track. When compared to Fig. 6, it is clear that the 2.5 fold increase in velocity does not comparably reduce the minimum source strength detectable.

All of the preceding effects serve to reduce the performance obtained with the current instrument. To balance this, the effective area of the current imager is 1.14 m² whereas the instrument under construction will only be 0.91 m² and its thickness will only be 4.2 cm versus 10 cm for the current detector. One would expect that the reduction in area will marginally affect the performance since the background will also be reduced by shielding. The thinner crystals will clearly impact the performance, but only at the highest energies.

Summary

We have taken a series of data sets that have allowed us to determine the sensitivity of our prototype large area imager search instrument. With these data sets we have generated receiver operator curves that show the instrument is sensitive to sources as weak as 1 mCi to distances in excess of 50 m. With dual sided imaging the prototype under construction should easily be able to sweep a path over 100-m wide for sources of this intensity.

References

1. K.P. Ziock, W.H. Goldstein, *The Lost Source, Varying Backgrounds and Why Bigger May not be Better*, in *Unattended Radiation Sensor Systems for Remote Applications*, J. Trombka, D. Spears, P. Solomon ed., AIP Conf. Proc. 632, (American Institute of Physics, Melville, NY, 2002) 60-70.

2. K.P. Ziock, W.W. Craig, L. Fabris, R.C. Lanza, S. Gallagher, B.K.P. Horn, N.W. Madden, *Large Area Imaging Detector for Long-Range, Passive Detection of Fissile Material*, IEEE Trans. Nuclear Science **51**, 2238-2244, 2004.
3. *Source-Search Sensitivity of a Large-Area, Coded-Aperture, Gamma-Ray Imager*, K.P. Ziock, W.W. Craig, L. Fabris, R. C. Lanza,, S. Gallagher, B.K.P. Horn, N.W. Madden, to be published in IEEE Trans. Nuclear Science.